

Exhibit 2

U.S. Patent No. 7,519,814 vs. HPE

VirtaMove asserts that Defendant Hewlett Packard Enterprise Company (“Defendant” or “HPE”) infringes the following claims (collectively, “Asserted Claims”): U.S. Patent No. 7,519,814 (“the ’814 patent”), claims 1, 2, 6, 9, and 10.

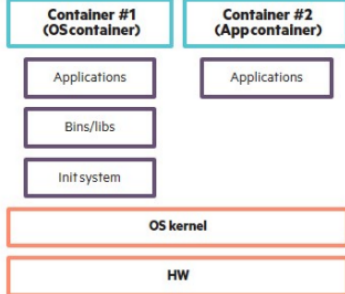
Accused Instrumentalities: HPE products and services using secure containerized applications, including without limitation HPE’s Ezmeral Runtime Enterprise (including without limitation both Ezmeral Runtime Enterprise and Ezmeral Runtime Enterprise Essentials, in each case including when marketed, sold, and/or licensed as part of or associated with HPE’s GreenLake branding, e.g. “HPE GreenLake for containers” which “is built on HPE Ezmeral Container Platform”), and all versions and variations thereof since the issuance of the asserted patent.

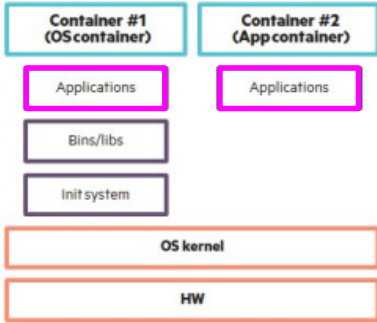
Each Accused Instrumentality infringes the claims in substantially the same way, and the evidence shown in this chart is similarly applicable to each Accused Instrumentality. Each claim limitation is literally infringed by each Accused Instrumentality. However, to the extent any claim limitation is not met literally, it is nonetheless met under the doctrine of equivalents because the differences between the claim limitation and each Accused Instrumentality would be insubstantial, and each Accused Instrumentality performs substantially the same function, in substantially the same way, to achieve the same result as the claimed invention. Notably, Defendant has not yet articulated which, if any, particular claim limitations it believes are not met by the Accused Instrumentalities.

Claim 1


Claim 1	Accused Instrumentalities
<p>[1pre] 1. In a system having a plurality of servers with operating systems that differ, operating in disparate computing environments, wherein each server includes a processor and an operating system including a kernel a set of associated local system files compatible with the processor, a method of providing at least some of the servers in the system with secure, executable, applications related to a service, wherein the applications are executed in a secure environment, wherein the applications each include an object executable by at least some of the different operating</p>	<p>To the extent the preamble is limiting, HPE and/or its customer practices, through the Accused Instrumentalities, in a system having a plurality of servers with operating systems that differ, operating in disparate computing environments, wherein each server includes a processor and an operating system including a kernel a set of associated local system files compatible with the processor, a method of providing at least some of the servers in the system with secure, executable, applications related to a service, wherein the applications are executed in a secure environment, wherein the applications each include an object executable by at least some of the different operating systems for performing a task related to the service, as claimed.</p> <p>For example, HPE Ezmeral Runtime Enterprise runs on individual servers, including HPE Synergy and HPE ProLiant servers, each of which runs an independent operating system, including for example RHEL or SLES running either on bare metal, through an on-premises virtualized infrastructure, through one or more cloud services, or through any other supported deployment. In an exemplary deployment, two or more servers use different operating systems. The servers operate in disparate computing environments, including because each server is a stand-alone computer and/or</p>

Claim 1	Accused Instrumentalities
<p>systems for performing a task related to the service, the method comprising:</p>	<p>each server is unrelated to the other servers due to having independent hardware and, in some instances, independent software.</p> <p>HPE requires that each server includes a processor with one or more cores available to the OS kernel. HPE further requires each server to have a supported operating system (SLES or RHEL/CentOS), which includes a kernel and associated local system files, including for example libraries such as libc/glibc, configuration files, etc. In the infringing system, at least two servers have different operating systems, for example SLES and RHEL/CentOS, or for another example different versions of SLES and/or RHEL/CentOS.</p> <p>In at least some instances, HPE directly owns, operates, controls, and/or benefits from the claimed system and/or method. In other instances, HPE's customer makes and uses the system and/or method either by following HPE's direction and control, including HPE's documentation, or automatically through the ordinary and expected operation of HPE's software, or a combination thereof.</p> <p><i>See claim limitations below.</i></p> <p><i>See also, e.g.:</i></p> <p>HPE Ezmeral Runtime Enterprise is an enterprise-grade container orchestration platform that is designed for the containerization of both cloud-native and non-cloud-native monolithic applications with persistent data. It deploys 100% open-source Kubernetes for orchestration, provides a state-of-the-art file system and data fabric for persistent container storage, and provides enterprises with the ability to deploy non-cloud-native AI and Analytics workloads in containers. Enterprises can now easily extend the agility and efficiency benefits of containers to more of their enterprise applications—running on either bare-metal or virtualized infrastructure, on-premises, in multiple clouds, or at the edge.</p> <p>https://www.hpe.com/psnow/doc/a50004264enw.pdf?jumpid=in_pdp-psnow-qs</p> <p>The offering formerly known as the HPE Ezmeral Container Platform is really focused on a lot more than just containers, and it provides businesses with more than just container orchestration software. The name change to HPE Ezmeral Runtime Enterprise reflects the fact that this is not just a solution for container platform orchestration. This platform offers an incredible wealth of capabilities and features you can use to modernize, deploy, monitor, and manage your applications.</p> <p>https://community.hpe.com/t5/hpe-ezmeral-uncut/hpe-ezmeral-container-platform-is-now-hpe-ezmeral-runtime/ba-p/7151720</p> <p>▪ OS agnostic – With an application and all its necessary files bundled into one unit – minus an operating system – the container can run on different operating systems, hardware, networks, storage systems and security policies. This means that any environment is compatible, so developers don't need to re-write applications for different servers.</p> <p>https://www.hpe.com/us/en/what-is/caas.html</p>

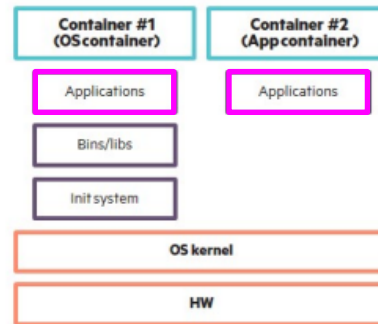
Claim 1	Accused Instrumentalities																
	<p>With HPE Ezmeral Runtime Enterprise running on HPE Synergy or HPE ProLiant servers, enterprises can extend the agility and efficiency benefits of containers to more of their enterprise applications—running on either bare metal or virtualized infrastructure, either on-premises, in multiple public clouds, or at the edge.</p> <p>https://www.hpe.com/psnow/doc/a50003599enw</p> <p>HPE Ezmeral Runtime Enterprise and HPE Ezmeral ML Ops Capabilities Matrix</p> <table><tr><th></th><th>HPE Ezmeral Runtime Enterprise Essentials</th><th>HPE Ezmeral Runtime Enterprise</th><th>HPE Ezmeral MLOps</th></tr><tr><td>Operating Systems (OS)</td><td>Yes</td><td>Yes</td><td>Yes</td></tr><tr><td>RHEL OS</td><td>Yes</td><td>Yes</td><td>Yes</td></tr><tr><td>SLES OS</td><td>Yes</td><td>Yes</td><td>Yes</td></tr></table> <p>https://www.hpe.com/psnow/doc/a50004264enw.pdf?jumpid=in_pdp-psnow-qs</p> <p>Standard Features</p> <ul style="list-style-type: none">Leverages portability of containers to run on any infrastructure (HPE or non-HPE) and any public cloud <p>https://www.hpe.com/psnow/doc/a50004264enw.pdf?jumpid=in_pdp-psnow-qs</p> <p>Two Linux containers on a single system</p>  <pre>graph TD subgraph Containers direction TB subgraph C1 [Container #1 OS container] A1[Applications] B1[Bins/libs] I1[Init system] end subgraph C2 [Container #2 App container] A2[Applications] end end C1 --- C2 C1 --- C2 --- OS[OS kernel] OS --- HW[HW]</pre>		HPE Ezmeral Runtime Enterprise Essentials	HPE Ezmeral Runtime Enterprise	HPE Ezmeral MLOps	Operating Systems (OS)	Yes	Yes	Yes	RHEL OS	Yes	Yes	Yes	SLES OS	Yes	Yes	Yes
	HPE Ezmeral Runtime Enterprise Essentials	HPE Ezmeral Runtime Enterprise	HPE Ezmeral MLOps														
Operating Systems (OS)	Yes	Yes	Yes														
RHEL OS	Yes	Yes	Yes														
SLES OS	Yes	Yes	Yes														

Claim 1	Accused Instrumentalities
	<p>https://h50146.www5.hpe.com/products/software/oe/linux/mainstream/support/whitepaper/pdfs/4AA6-2761ENW.pdf</p> <p>Each license allows the customer to deploy the HPE Ezmeral Container Platform on one Core and 2 terabytes of Storage Capacity. The customer must purchase more licenses if they exceed the allowable amount of Cores or Storage Capacity. As used in this Agreement, Core means a part of a CPU that executes a single stream of compiled instruction code. Each physical processor contains smaller processing units called physical CPU cores. Some processors have two cores, some four, some eight, and so on. Core capacity represents the total number of cores available within a given system. The number of cores is counted as the number of logical cores presented to the product guest OS. For licensing purposes, the number of cores on a given Ezmeral Container Platform host is the number of unique cores available to the kernel in the OS on which the Ezmeral Container Platform software is directly installed, regardless of the number of threads in each core. It equals the product of Core(s) per socket and Socket(s), as shown in the output of https://docs.ezmeral.hpe.com/runtime-enterprise/56/home/about-hpe-ezmeral-container-pl/GEN_End_User_Software_Agreement.html</p> <p>Kernel mode refers to the processor mode that enables software to have full and unrestricted access to the system and its resources. The OS kernel and kernel drivers, such as the file system driver, are loaded into protected memory space and operate in this highly privileged kernel mode.</p> <p>https://www.techtarget.com/searchdatacenter/definition/kernel</p> <p>Instead of using a hypervisor to manage VMs, the figure shows how containers isolate applications into separate environments (containers) that include processor, memory, and networking resources as part of the container itself. This environment provides OS-level virtualization. Containers have their own root; and, users and processes do not perform operations outside of the container environment. The host OS kernel manages container workloads directly, which reduces the overhead involved with managing system resources. This improves efficiency and therefore, improves performance.</p> <p>Two Linux containers on a single system</p>  <pre> graph TD subgraph Containers direction TB C1[Container #1 (OS container)] C2[Container #2 (App container)] end subgraph Layers direction TB A1[Applications] B1[Bins/libs] I1[Init system] K[OS kernel] H[HW] end C1 --- A1 C1 --- B1 C1 --- I1 C2 --- A2[Applications] C2 --- B2[Bins/libs] C2 --- I2[Init system] A1 --- K A2 --- K B1 --- K B2 --- K I1 --- K I2 --- K K --- H </pre>

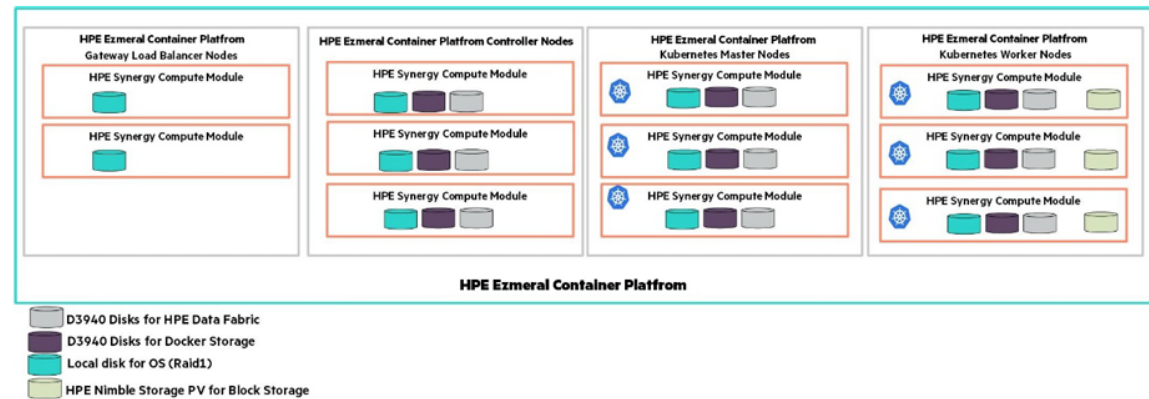
Claim 1	Accused Instrumentalities
	<p data-bbox="674 199 1965 264">https://h50146.www5.hpe.com/products/software/oe/linux/mainstream/support/whitepaper/pdfs/4AA6-2761ENW.pdf</p> <p data-bbox="674 310 1627 451">Docker container: A <i>Docker container</i> is a lightweight, standalone, executable software package that runs specific services. This software package includes code, runtime, system libraries, configurations, etc. that run as an isolated process in user space. A Docker container is typically used to deploy scalable and repeatable <i>microservices</i>. HPE Ezmeral Runtime Enterprise contains innovations around storage, networking, and security to utilize Docker containers as lightweight virtual machines to run Big Data and analytics applications.</p> <p data-bbox="674 464 1965 529">https://support.hpe.com/hpesc/public/docDisplay?docId=a00ecp54hen_us&docLocale=en_US&page=home/about-hpe-ezmeral-container-pl/GEN_Definitions.html</p> <h2 data-bbox="674 553 1612 610">Controller, Gateway, and Worker Hosts</h2> <p data-bbox="674 662 1927 789">A host is either a physical server or a virtual server, located on your premises or in a public cloud, that is available to HPE Ezmeral Runtime Enterprise. The term host and node are often used interchangeably. Nodes are hosts that are part of a cluster.</p> <p data-bbox="674 837 1938 964">You must have a supported operating system installed on hosts before they can be used in HPE Ezmeral Runtime Enterprise. Hosts have different requirements depending on their functions. See Host Requirements.</p> <p data-bbox="674 997 1596 1062">https://docs.ezmeral.hpe.com/runtime-enterprise/56/reference/universal-concepts/Controller_Gateway_and_Worker_Hosts.html</p>

Claim 1	Accused Instrumentalities
	<p>✓ Kubernetes Cluster Nodes </p> <p>A deployment of HPE Ezmeral Runtime Enterprise can include multiple Kubernetes clusters. A host that is part of a Kubernetes cluster is referred to in Kubernetes as a node.</p> <p>Each Kubernetes cluster has its own control plane, consisting of at least one control plane node. The Kubernetes control plane is separate from the Platform Control Plane. A high-availability Kubernetes cluster has multiple control plane nodes, as described in High Availability.</p> <p>Kubernetes clusters contain worker nodes that run the containers and pods that process jobs in HPE Ezmeral Runtime Enterprise.</p> <p>For more information about hosts and Kubernetes clusters, see Controller, Gateway, and Worker Hosts.</p> <p>https://docs.ezmeral.hpe.com/runtime-enterprise/56/reference/kubernetes/Kubernetes_Physical_Architecture.html#v52_k8s-kubernetes-physical-architecture_k8s-cluster-architecture</p>
<p>[1a] storing in memory accessible to at least some of the servers a plurality of secure containers of application software, each container comprising one or more of the executable applications and a set of associated system files required to execute the one or more applications, for use with a local kernel residing permanently on one of the servers;</p>	<p>The method practiced by HPE and/or its customer through the Accused Instrumentalities includes a step of storing in memory accessible to at least some of the servers a plurality of secure containers of application software, each container comprising one or more of the executable applications and a set of associated system files required to execute the one or more applications, for use with a local kernel residing permanently on one of the servers.</p> <p>For example, HPE Ezmeral Runtime Enterprise stores application containers, sometimes called Docker containers, container images, Kubernetes containers, or Kubernetes pods, in persistent storage available to each node running the application. The terms “node” and “host” are both used to refer to the claimed server. The container might be in a format defined by the Open Container Initiative. This storage may be physically attached to the server or connected through any supported interconnect, including over a network. In addition to the application software, each container includes associated system files, including a Linux user space required to execute the application, for example libc/glibc and other shared libraries, configuration files, etc. necessary for the application. For example, the container includes a base OS image, provided by HPE or by a third party, such as a CentOS, RHEL,</p>

Claim 1	Accused Instrumentalities
	<p>or Ubuntu base image. The container is compatible with the host kernel, for example because the container libraries are linked against the Linux kernel, and the supported host operating systems also use the Linux kernel, which has a stable binary interface.</p> <p>The containers are secure containers as claimed. For example, the data within an individual container is insulated from the effects of other containers except to the extent the container is specifically configured to allow other containers to modify its data, for example using a shared volume.</p> <p><i>See, e.g.:</i></p> <p>Pod: For Kubernetes, a <i>pod</i> is a group of containers deployed on a single host.</p> <p>Data Fabric cluster: This is a Kubernetes cluster that is used for HPE Ezmeral Data Fabric storage. A Data Fabric cluster is a Custom Resource in Kubernetes that is supported by operators in HPE Ezmeral Runtime Enterprise.</p> <p>Data Fabric CR: This typically refers to the Custom Resource specification for a Data Fabric cluster that is supported by an HPE Ezmeral Runtime Enterprise <i>dataplat</i>form operator. It specifies each type of pod that the cluster would comprise. The per-pod specification may include CPU, memory, disk, and port requirements. Together with node labels and annotations, the Data Fabric CR influences the placement and scheduling of cluster pods by Kubernetes. HPE Ezmeral Runtime Enterprise creates and applies the Data Fabric CR when creating the first Data Fabric cluster. The Data Fabric CR may be subsequently patched/modified when expanding the cluster, or by a user with suitable privileges.</p> <p>Docker container: A <i>Docker container</i> is a lightweight, standalone, executable software package that runs specific services. This software package includes code, runtime, system libraries, configurations, etc. that run as an isolated process in user space. A Docker container is typically used to deploy scalable and repeatable <i>microservices</i>. HPE Ezmeral Runtime Enterprise contains innovations around storage, networking, and security to utilize Docker containers as lightweight virtual machines to run Big Data and analytics applications.</p> <p>https://support.hpe.com/hpesc/public/docDisplay?docId=a00ecp54hen_us&docLocale=en_US&page=home/about-hpe-ezmeral-container-pl/GEN_Definitions.html</p>

Claim 1**Accused Instrumentalities****Two Linux containers on a single system**

<https://h50146.www5.hpe.com/products/software/oe/linux/mainstream/support/whitepaper/pdfs/4AA6-2761ENW.pdf>

**Storage**

The HPE Synergy D3940 Storage Module provides solid state disks and hard disk drives to local systems where it is consumed by HPE Ezmeral Data Fabric, and optionally by compute nodes as boot devices. HPE Nimble Storage dynamically provides Persistent Volume (PV) for containers using Dynamic Volume Provisioner which is integrated with the HPE CSI Driver.

<https://www.hpe.com/psnow/doc/a50002075enw>

HPE Nimble Storage

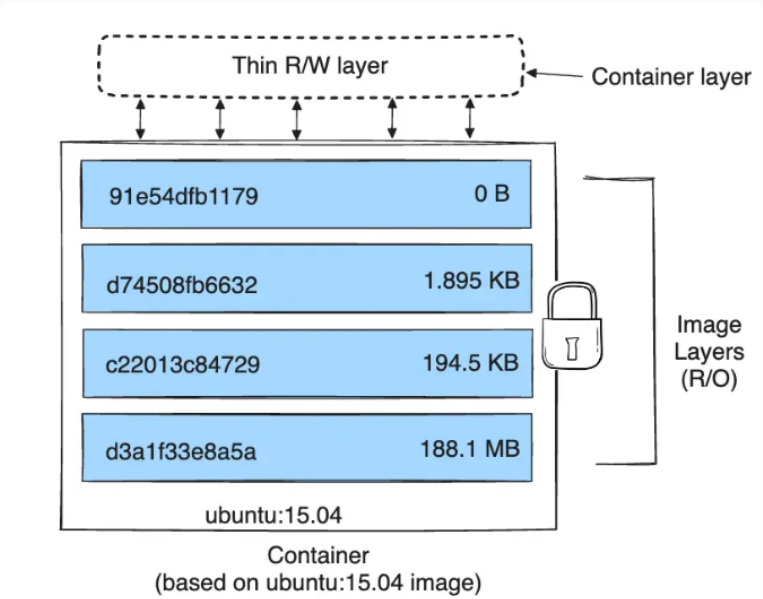
HPE Nimble Storage AF40 is used to provide persistent, block storage in this solution. The HPE Nimble Storage array for Docker data provides the storage volume to host the repository, to store container images, and also provides persistent volume for applications.

Claim 1	Accused Instrumentalities
	<p>https://www.hpe.com/psnow/doc/a50002075enw</p> <h2 data-bbox="709 321 1020 370">Container images</h2> <p data-bbox="709 397 1304 524">A container image is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <p>https://kubernetes.io/docs/concepts/containers/</p> <p data-bbox="693 605 1808 706">An application container is a stand-alone, all-in-one package for a software application. Containers include the application binaries, plus the software dependencies and the hardware requirements needed to run, all wrapped up into an independent, self-contained unit.</p> <p>https://developer.hpe.com/blog/kubernetes-application-containers-managing-containers-and-cluster-resour/</p> <p data-bbox="678 813 1461 891">Because each application container creates an isolated environment for its application, the resources allocated to it are the entire machine. Other copies of the same container are "unaware" of each other.</p> <p>https://developer.hpe.com/blog/kubernetes-application-containers-managing-containers-and-cluster-resour/</p> <ul data-bbox="688 992 1963 1076" style="list-style-type: none"> ▪ OS agnostic – With an application and all its necessary files bundled into one unit – minus an operating system – the container can run on different operating systems, hardware, networks, storage systems and security policies. This means that any environment is compatible, so developers don't need to re-write applications for different servers. <p>https://www.hpe.com/us/en/what-is/caas.html</p> <p data-bbox="688 1206 1451 1230">6. Do Docker containers package up the entire OS and make it easier to deploy?</p> <p data-bbox="688 1268 1709 1352">Docker containers do not package up the OS. They package up the applications with everything that the application needs to run. The engine is installed on top of the OS running on a host. Containers share the OS kernel allowing a single host to run multiple containers.</p> <p>https://www.docker.com/blog/the-10-most-common-questions-it-admins-ask-about-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Kubernetes namespaces have the following uses:</p> <ul style="list-style-type: none">• Isolation: Teams, projects, and customers exist in their own environment within a cluster, and do not impact each other's work. <p>https://support.hpe.com/hpesc/public/docDisplay?docId=a00ecp55hen_us&docLocale=en_US&page=reference/universal-concepts/Namespaces.html</p> <p>Using containers isolates software and allows it to work independently across different operating systems, hardware, networks, storage systems, and security policies. It allows the container-based application to transition seamlessly through development, testing, and production environments. Because an operating system is not packed into the container, each container uses minimal computing resources, making it light and easy to install.</p> <p>https://www.hpe.com/us/en/what-is/containers.html</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="682 212 1312 277">About storage drivers</h2> <p data-bbox="682 326 1913 448">To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2 data-bbox="682 518 1604 574">Storage drivers versus Docker volumes</h2> <p data-bbox="682 612 1953 878">Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p data-bbox="682 927 1944 1052">Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the <a data-bbox="1381 976 1575 1003" href="#">volumes section to learn how to use volumes to persist data and improve performance.</p> <p data-bbox="674 1084 1266 1117"><a data-bbox="674 1084 1266 1117" href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</p>

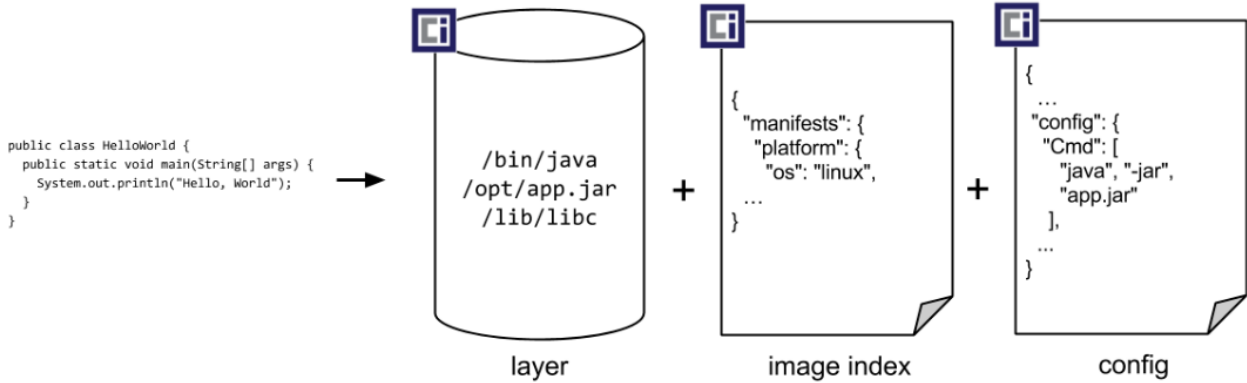
Claim 1	Accused Instrumentalities
	<h2 data-bbox="695 207 1121 261">Images and layers</h2> <p data-bbox="695 302 1860 375">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="695 418 1946 782"> # syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py </pre> <p data-bbox="695 824 1940 1133">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="674 1154 1266 1187">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the <code>\$HOME/.cache</code> directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the Best practices for writing Dockerfiles and use multi-stage builds sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p>  <p>The diagram illustrates the layer architecture of a container. At the bottom is a box labeled 'Container (based on ubuntu:15.04 image)'. Inside this box is a stack of four blue rectangular layers representing the 'Image Layers (R/O)'. From top to bottom, the layers are labeled with their IDs and sizes: '91e54dfb1179' (0 B), 'd74508fb6632' (1.895 KB), 'c22013c84729' (194.5 KB), and 'd3a1f33e8a5a' (188.1 MB). A padlock icon is positioned to the right of the stack, indicating that these layers are read-only. Above the stack is a dashed rectangular box labeled 'Thin R/W layer'. An arrow points from the text 'Container layer' to this dashed box. Vertical double-headed arrows connect the top of the 'Thin R/W layer' to the top of each of the four image layers, showing the relationship between the container's writable layer and the underlying image layers.</p> <p>https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="690 224 959 282">Volumes</h2> <p data-bbox="690 337 1944 467">Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While bind mounts are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p> <p data-bbox="672 492 1346 521">https://kubernetes.io/docs/concepts/storage/volumes/</p> <h2 data-bbox="697 573 1264 621">Container environment</h2> <p data-bbox="697 662 1514 727">The Kubernetes Container environment provides several important resources to Containers:</p> <ul data-bbox="735 764 1488 922" style="list-style-type: none">• A filesystem, which is a combination of an image and one or more volumes.• Information about the Container itself.• Information about other objects in the cluster. <p data-bbox="672 959 1566 989">https://kubernetes.io/docs/concepts/containers/container-environment/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="695 215 915 277">Images</h2> <p data-bbox="695 310 1562 461">A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p> <p data-bbox="695 500 1568 570">You typically create a container image of your application and push it to a registry before referring to it in a Pod.</p> <p data-bbox="674 597 1367 630">https://kubernetes.io/docs/concepts/containers/images/</p> <h2 data-bbox="695 675 957 737">Volumes</h2> <p data-bbox="695 773 1568 1214">On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers. One problem occurs when a container crashes or is stopped. Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a Pod and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes volume abstraction solves both of these problems. Familiarity with Pods is suggested.</p> <p data-bbox="674 1242 1346 1274">https://kubernetes.io/docs/concepts/storage/volumes/</p>


Claim 1	Accused Instrumentalities
	<div data-bbox="709 219 1337 277"><h2>Open Container Initiative</h2><hr/></div> <div data-bbox="709 339 1222 386"><h3>Image Format Specification</h3><hr/></div> <div data-bbox="709 433 1944 506"><p>This specification defines an OCI Image, consisting of an image manifest, an image index (optional), a set of filesystem layers, and a configuration.</p></div> <div data-bbox="709 542 1944 615"><p>The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p></div> <div data-bbox="669 644 1520 714"><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p></div>

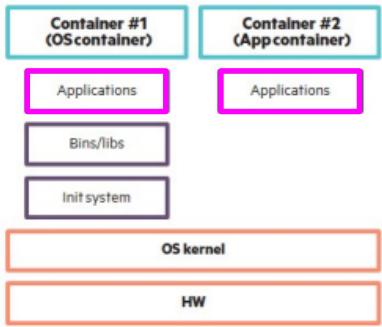
Claim 1	Accused Instrumentalities
	<p>Overview</p> <p>At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more filesystem layer changeset archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p>  <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="690 207 1335 266">OCI Image Configuration</h2> <p data-bbox="690 321 1955 483">An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in Layers.</p> <p data-bbox="690 521 1696 553">This section defines the <code>application/vnd.oci.image.config.v1+json</code> media type.</p> <p data-bbox="672 586 1543 656">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 1	Accused Instrumentalities
	<p data-bbox="701 215 789 248">Layer</p> <ul data-bbox="730 289 1957 638" style="list-style-type: none"> • Image filesystems are composed of <i>layers</i>. • Each layer represents a set of filesystem changes in a tar-based layer format, recording files to be added, changed, or deleted relative to its parent layer. • Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer. • Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem. <p data-bbox="701 686 898 719">Image JSON</p> <ul data-bbox="730 760 1957 1109" style="list-style-type: none"> • Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes. • The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers. • This JSON is considered to be immutable, because changing it would change the computed ImageID. • Changing it means creating a new derived image, instead of changing the existing image. <p data-bbox="674 1133 1543 1206"> https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md </p>

Claim 1	Accused Instrumentalities
	<ul style="list-style-type: none"> • rootfs <i>object</i>, REQUIRED <p>The rootfs key references the layer content addresses used by the image. This makes the image config hash depend on the filesystem hash.</p> <ul style="list-style-type: none"> ◦ type <i>string</i>, REQUIRED <p>MUST be set to <code>layers</code>. Implementations MUST generate an error if they encounter a unknown value while verifying or unpacking an image.</p> <ul style="list-style-type: none"> ◦ diff_ids <i>array of strings</i>, REQUIRED <p>An array of layer content hashes (<code>DiffIDs</code>), in order from first to last.</p> <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p> <h2>Controller, Gateway, and Worker Hosts</h2> <p>A host is either a physical server or a virtual server, located on your premises or in a public cloud, that is available to HPE Ezmeral Runtime Enterprise. The term host and node are often used interchangeably. Nodes are hosts that are part of a cluster.</p> <p>You must have a supported operating system installed on hosts before they can be used in HPE Ezmeral Runtime Enterprise. Hosts have different requirements depending on their functions. See Host Requirements.</p> <p>https://docs.ezmeral.hpe.com/runtime-enterprise/56/reference/universal-concepts/Controller_Gateway_and_Worker_Hosts.html</p>

Claim 1	Accused Instrumentalities
	<p data-bbox="695 199 1178 232">✓ Kubernetes Cluster Nodes </p> <p data-bbox="695 284 1955 362">A deployment of HPE Ezmeral Runtime Enterprise can include multiple Kubernetes clusters. A host that is part of a Kubernetes cluster is referred to in Kubernetes as a node.</p> <p data-bbox="695 410 1913 540">Each Kubernetes cluster has its own control plane, consisting of at least one control plane node. The Kubernetes control plane is separate from the Platform Control Plane. A high-availability Kubernetes cluster has multiple control plane nodes, as described in High Availability.</p> <p data-bbox="695 586 1913 667">Kubernetes clusters contain worker nodes that run the containers and pods that process jobs in HPE Ezmeral Runtime Enterprise.</p> <p data-bbox="695 712 1934 745">For more information about hosts and Kubernetes clusters, see Controller, Gateway, and Worker Hosts.</p> <p data-bbox="674 781 1913 886">https://docs.ezmeral.hpe.com/runtime-enterprise/56/reference/kubernetes/Kubernetes_Physical_Architecture.html#v52_k8s-kubernetes-physical-architecture__k8s-cluster-architecture</p>

Claim 1	Accused Instrumentalities
<p>[1b] wherein the set of associated system files are compatible with a local kernel of at least some of the plurality of different operating systems,</p>	<p>In the method practiced by HPE and/or its customer through the Accused Instrumentalities, the set of associated system files are compatible with a local kernel of at least some of the plurality of different operating systems.</p> <p>The associated system files in the container are compatible with the host kernel, for example because they are linked against the Linux kernel and the supported host operating systems also use the Linux kernel, which has a stable binary interface.</p> <p>See discussion in element [1a] above.</p> <p>See, e.g.:</p> <p>Docker container: A <i>Docker container</i> is a lightweight, standalone, executable software package that runs specific services. This software package includes code, runtime, system libraries, configurations, etc. that run as an isolated process in user space. A Docker container is typically used to deploy scalable and repeatable <i>microservices</i>. HPE Ezmeral Runtime Enterprise contains innovations around storage, networking, and security to utilize Docker containers as lightweight virtual machines to run Big Data and analytics applications.</p> <p>https://support.hpe.com/hpesc/public/docDisplay?docId=a00ecp54hen_us&docLocale=en_US&page=home/about-hpe-ezmeral-container-pl/GEN_Definitions.html</p> <p>▪ OS agnostic – With an application and all its necessary files bundled into one unit – minus an operating system – the container can run on different operating systems, hardware, networks, storage systems and security policies. This means that any environment is compatible, so developers don't need to re-write applications for different servers.</p> <p>https://www.hpe.com/us/en/what-is/caas.html</p> <p>Two Linux containers on a single system</p>  <pre> graph TD subgraph Containers direction LR C1[Container #1
(OScontainer)] C2[Container #2
(Appcontainer)] end A1[Applications] --- C1 A2[Applications] --- C2 B1[Bins/libs] --- C1 I1[Init system] --- C1 C1 --- OK[OS kernel] C2 --- OK OK --- HW[HW] </pre>

Claim 1	Accused Instrumentalities
	https://h50146.www5.hpe.com/products/software/oe/linux/mainstream/support/whitepaper/pdfs/4AA6-2761ENW.pdf

Claim 1	Accused Instrumentalities
<p>[1c] the containers of application software excluding a kernel,</p>	<p>In the method practiced by HPE and/or its customer through the Accused Instrumentalities, the containers of application software exclude a kernel.</p> <p><i>See, e.g.:</i></p> <p>Docker container: A <i>Docker container</i> is a lightweight, standalone, executable software package that runs specific services. This software package includes code, runtime, system libraries, configurations, etc. that run as an isolated process in user space. A Docker container is typically used to deploy scalable and repeatable <i>microservices</i>. HPE Ezmeral Runtime Enterprise contains innovations around storage, networking, and security to utilize Docker containers as lightweight virtual machines to run Big Data and analytics applications.</p> <p>https://support.hpe.com/hpesc/public/docDisplay?docId=a00ecp54hen_us&docLocale=en_US&page=home/about-hpe-ezmeral-container-pl/GEN_Definitions.html</p> <p>6. Do Docker containers package up the entire OS and make it easier to deploy?</p> <p>Docker containers do not package up the OS. They package up the applications with everything that the application needs to run. The engine is installed on top of the OS running on a host. Containers share the OS kernel allowing a single host to run multiple containers.</p> <p>https://www.docker.com/blog/the-10-most-common-questions-it-admins-ask-about-docker/</p> <p>Containers and VMs perform somewhat similar functions in that they provide virtualized environments in which software applications can run separately from the rest of the system. But these technologies are very different and are used in different situations. Each virtual machine runs both an OS and the application, while containers share a single OS via a kernel, making them more lightweight and portable.</p> <p>https://www.hpe.com/us/en/what-is/containers.html</p>

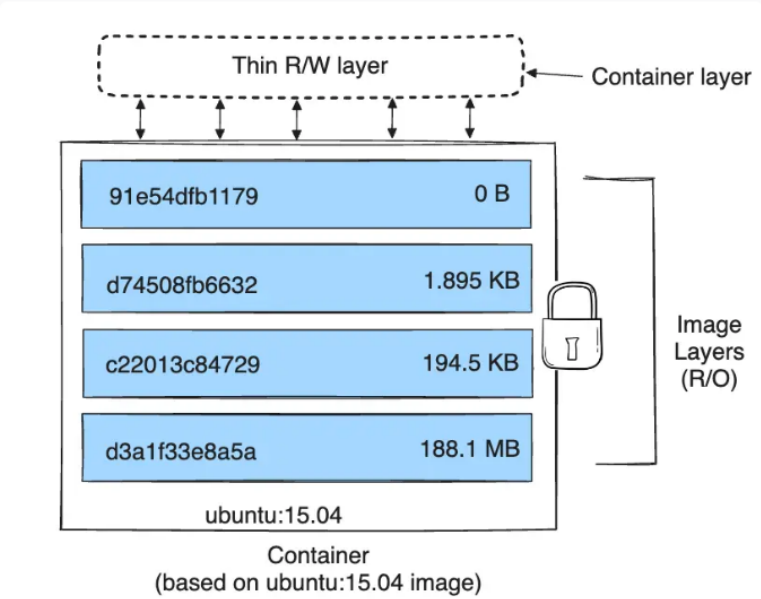
Claim 1	Accused Instrumentalities
<p>[1d] wherein some or all of the associated system files within a container stored in memory are utilized in place of the associated local system files that remain resident on the server,</p>	<p>In the method practiced by HPE and/or its customer through the Accused Instrumentalities, some or all of the associated system files within a container stored in memory are utilized in place of the associated local system files that remain resident on the server.</p> <p>For example, each container will utilize its own associated system files, including libraries such as libc/glibc and configuration files, not the corresponding associated local system files (<i>e.g.</i>, libraries and configuration files of the host OS). As described above and below, in the Accused Instrumentalities the associated system files provide at least some of the same functionalities as the associated local system files. The host/node's associated local system files remain resident on the host/node, for example for use by system processes or applications outside the container environment.</p> <p><i>See, e.g.:</i></p> <p>Docker container: A <i>Docker container</i> is a lightweight, standalone, executable software package that runs specific services. This software package includes code, runtime, system libraries, configurations, etc. that run as an isolated process in user space. A Docker container is typically used to deploy scalable and repeatable <i>microservices</i>. HPE Ezmeral Runtime Enterprise contains innovations around storage, networking, and security to utilize Docker containers as lightweight virtual machines to run Big Data and analytics applications.</p> <p>https://support.hpe.com/hpesc/public/docDisplay?docId=a00ecp54hen_us&docLocale=en_US&page=home/about-hpe-ezmeral-container-pl/GEN_Definitions.html</p> <p>An application container is a stand-alone, all-in-one package for a software application. Containers include the application binaries, plus the software dependencies and the hardware requirements needed to run, all wrapped up into an independent, self-contained unit.</p> <p>https://developer.hpe.com/blog/kubernetes-application-containers-managing-containers-and-cluster-resour/</p> <p>▪ OS agnostic – With an application and all its necessary files bundled into one unit – minus an operating system – the container can run on different operating systems, hardware, networks, storage systems and security policies. This means that any environment is compatible, so developers don't need to re-write applications for different servers.</p> <p>https://www.hpe.com/us/en/what-is/caas.html</p>

Claim 1	Accused Instrumentalities
<p>[1e] wherein said associated system files utilized in place of the associated local system files are copies or modified copies of the associated local system files that remain resident on the server,</p>	<p>In the method practiced by HPE and/or its customer through the Accused Instrumentalities, said associated system files utilized in place of the associated local system files are copies or modified copies of the associated local system files that remain resident on the server.</p> <p>For example, in some cases the host OS and container will use one or more identical system files, for example when both the host and the container incorporate the same Linux distribution version, or when both host and container use the same version of libc. In the case where the associated system files are identical to the associated local system files, they are copies thereof. In other cases modified copies are used instead, for example when different versions of the same library, or configuration files with different parameters, are used by the host and container.</p> <p><i>See, e.g.:</i></p> <p>Docker container: A <i>Docker container</i> is a lightweight, standalone, executable software package that runs specific services. This software package includes code, runtime, system libraries, configurations, etc. that run as an isolated process in user space. A Docker container is typically used to deploy scalable and repeatable <i>microservices</i>. HPE Ezmeral Runtime Enterprise contains innovations around storage, networking, and security to utilize Docker containers as lightweight virtual machines to run Big Data and analytics applications.</p> <p>https://support.hpe.com/hpesc/public/docDisplay?docId=a00ecp54hen_us&docLocale=en_US&page=home/about-hpe-ezmeral-container-pl/GEN_Definitions.html</p> <p>COPY and ADD : These commands copy files and directories from your local filesystem into the Docker image. They are often used to include your application code, configuration files, and dependencies.</p> <p>https://medium.com/@swalperen3008/what-is-dockerize-and-dockerize-your-project-a-step-by-step-guide-899c48a34df6</p> <h3>Container images</h3> <p>A <i>container image</i> is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <p>https://kubernetes.io/docs/concepts/containers/</p>

Claim 1	Accused Instrumentalities
<p>[1f] and wherein the application software cannot be shared between the plurality of secure containers of application software,</p>	<p>In the method practiced by HPE and/or its customer through the Accused Instrumentalities, the application software cannot be shared between the plurality of secure containers of application software.</p> <p>For example, each container has an isolated runtime environment that cannot be accessed by other containers, for example including a per-container writeable layer or other ephemeral per-container storage. For another example, when the plurality of secure containers each corresponds to a different container image, each container cannot access another container's image and therefore application software.</p> <p>See, e.g.:</p> <p>Docker container: A <i>Docker container</i> is a lightweight, standalone, executable software package that runs specific services. This software package includes code, runtime, system libraries, configurations, etc. that run as an isolated process in user space. A Docker container is typically used to deploy scalable and repeatable <i>microservices</i>. HPE Ezmeral Runtime Enterprise contains innovations around storage, networking, and security to utilize Docker containers as lightweight virtual machines to run Big Data and analytics applications.</p> <p>https://support.hpe.com/hpesc/public/docDisplay?docId=a00ecp54hen_us&docLocale=en_US&page=home/about-hpe-ezmeral-container-pl/GEN_Definitions.html</p> <p>Kubernetes namespaces have the following uses:</p> <ul style="list-style-type: none"> • Isolation: Teams, projects, and customers exist in their own environment within a cluster, and do not impact each other's work. <p>https://support.hpe.com/hpesc/public/docDisplay?docId=a00ecp55hen_us&docLocale=en_US&page=reference/universal-concepts/Namespaces.html</p> <p>Because each application container creates an isolated environment for its application, the resources allocated to it are the entire machine. Other copies of the same container are "unaware" of each other.</p> <p>https://developer.hpe.com/blog/kubernetes-application-containers-managing-containers-and-cluster-resour/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="682 212 1312 277">About storage drivers</h2> <p data-bbox="682 326 1913 448">To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2 data-bbox="682 518 1604 574">Storage drivers versus Docker volumes</h2> <p data-bbox="682 612 1953 878">Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p data-bbox="682 927 1944 1052">Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the <a data-bbox="1381 976 1575 1003" href="#">volumes section to learn how to use volumes to persist data and improve performance.</p> <p data-bbox="674 1084 1266 1117"><a data-bbox="674 1084 1266 1117" href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</p>

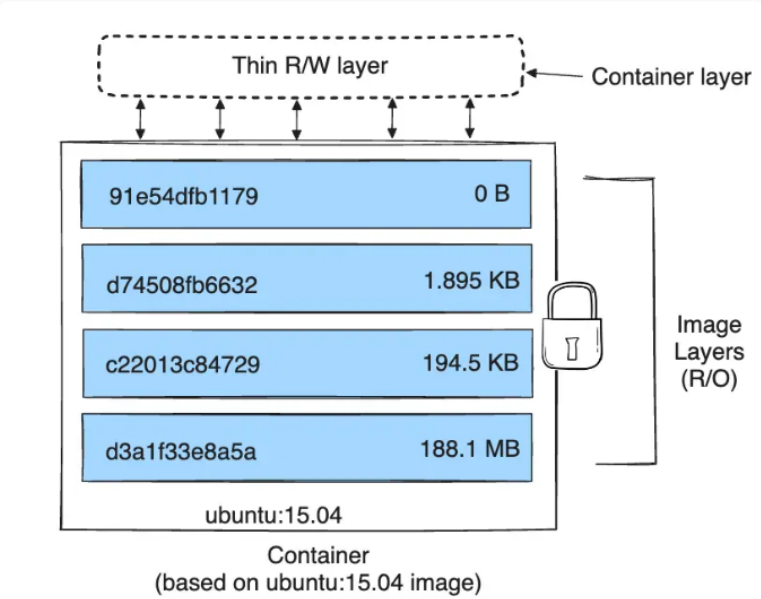
Claim 1	Accused Instrumentalities
	<h2 data-bbox="695 207 1121 261">Images and layers</h2> <p data-bbox="695 302 1860 375">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="695 418 1493 764"> # syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py </pre> <p data-bbox="695 824 1940 1130">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="674 1154 1268 1187">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the <code>\$HOME/.cache</code> directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the Best practices for writing Dockerfiles and use multi-stage builds sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p>  <p>The diagram illustrates the layer architecture of a container. At the bottom is a box labeled 'Container (based on ubuntu:15.04 image)'. Inside this box is a stack of four blue rectangular layers, each representing an image layer (R/O). The layers are labeled with their IDs and sizes: '91e54dfb1179' (0 B), 'd74508fb6632' (1.895 KB), 'c22013c84729' (194.5 KB), and 'd3a1f33e8a5a' (188.1 MB). A padlock icon is shown next to the stack, indicating that these layers are read-only. Above the stack is a dashed box labeled 'Thin R/W layer', which is identified as the 'Container layer' by an arrow. Double-headed vertical arrows connect the 'Thin R/W layer' to each of the four image layers below it, showing the relationship between the container's writable layer and the underlying image layers.</p> <p>https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
<p>[1g] and wherein each of the containers has a unique root file system that is different from an operating system's root file system.</p>	<p>In the method practiced by HPE and/or its customer through the Accused Instrumentalities, each of the containers has a unique root file system that is different from an operating system's root file system.</p> <p>For example, the container's root file system comprises the image layer(s), an ephemeral writeable layer (e.g., in Docker terminology the container layer), and optionally one or more volumes. This root file system is distinct and isolated from the host operating system's root file system.</p> <p><i>See, e.g.:</i></p> <p>Using containers isolates software and allows it to work independently across different operating systems, hardware, networks, storage systems, and security policies. It allows the container-based application to transition seamlessly through development, testing, and production environments. Because an operating system is not packed into the container, each container uses minimal computing resources, making it light and easy to install.</p> <p>https://www.hpe.com/us/en/what-is/containers.html</p> <p>Node storage: <i>Node storage</i> is storage space available for backing the <u>root file systems</u> of containers. Each HPE Ezmeral Runtime Enterprise host contributes node storage space that is used by the virtual nodes (Docker containers) assigned to that host. The Platform Administrator may optionally specify a quota limiting how much node storage a tenant's virtual nodes may consume.</p> <p>https://support.hpe.com/hpesc/public/docDisplay?docId=a00ecp54hen_us&docLocale=en_US&page=home/about-hpe-ezmeral-container-pl/GEN_Definitions.html</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="682 212 1312 277">About storage drivers</h2> <p data-bbox="682 326 1913 448">To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2 data-bbox="682 518 1604 574">Storage drivers versus Docker volumes</h2> <p data-bbox="682 612 1953 878">Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p data-bbox="682 927 1944 1052">Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the <a data-bbox="1381 976 1575 1003" href="#">volumes section to learn how to use volumes to persist data and improve performance.</p> <p data-bbox="674 1084 1266 1117"><a data-bbox="674 1084 1266 1117" href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="695 207 1121 261">Images and layers</h2> <p data-bbox="695 302 1860 375">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="695 418 1493 764"> # syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py </pre> <p data-bbox="695 824 1940 1130">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="674 1154 1266 1187">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the <code>\$HOME/.cache</code> directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the Best practices for writing Dockerfiles and use multi-stage builds sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p>  <p>The diagram illustrates the layer architecture of a container. At the bottom is a box labeled 'Container (based on ubuntu:15.04 image)'. Inside this box is a stack of four blue rectangular blocks representing image layers. From bottom to top, the layers are labeled with their IDs and sizes: 'd3a1f33e8a5a' (188.1 MB), 'c22013c84729' (194.5 KB), 'd74508fb6632' (1.895 KB), and '91e54dfb1179' (0 B). To the right of this stack is a bracket labeled 'Image Layers (R/O)' with a padlock icon, indicating they are read-only. Above the stack is a dashed box labeled 'Thin R/W layer'. An arrow points from the text 'Container layer' to this dashed box. Vertical double-headed arrows connect the top of the 'Thin R/W layer' to the top of each of the four image layers below it.</p> <p>https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="690 224 959 282">Volumes</h2> <p data-bbox="690 339 1944 467">Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While bind mounts are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p> <p data-bbox="672 492 1346 521">https://kubernetes.io/docs/concepts/storage/volumes/</p> <h2 data-bbox="697 574 1264 623">Container environment</h2> <p data-bbox="697 660 1514 725">The Kubernetes Container environment provides several important resources to Containers:</p> <ul data-bbox="735 764 1488 922" style="list-style-type: none">• A filesystem, which is a combination of an image and one or more volumes.• Information about the Container itself.• Information about other objects in the cluster. <p data-bbox="672 959 1566 989">https://kubernetes.io/docs/concepts/containers/container-environment/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="695 212 919 277">Images</h2> <p data-bbox="695 310 1562 461">A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p> <p data-bbox="695 498 1570 570">You typically create a container image of your application and push it to a registry before referring to it in a Pod.</p> <p data-bbox="674 597 1367 630">https://kubernetes.io/docs/concepts/containers/images/</p> <h2 data-bbox="695 672 961 737">Volumes</h2> <p data-bbox="695 769 1570 1214">On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers. One problem occurs when a container crashes or is stopped. Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a Pod and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes volume abstraction solves both of these problems. Familiarity with Pods is suggested.</p> <p data-bbox="674 1242 1346 1274">https://kubernetes.io/docs/concepts/storage/volumes/</p>

Claim 1	Accused Instrumentalities
	<div data-bbox="711 217 1337 277"><h2>Open Container Initiative</h2><hr/></div> <div data-bbox="711 339 1222 386"><h3>Image Format Specification</h3><hr/></div> <div data-bbox="711 433 1944 506"><p>This specification defines an OCI Image, consisting of an image manifest, an image index (optional), a set of filesystem layers, and a configuration.</p></div> <div data-bbox="711 542 1944 615"><p>The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p></div> <div data-bbox="669 644 1520 714"><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p></div>

Claim 1	Accused Instrumentalities
	<p data-bbox="688 212 869 250">Overview</p> <p data-bbox="688 306 1934 548">At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more filesystem layer changeset archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p> <div data-bbox="709 589 1948 971"> <pre data-bbox="709 727 1024 816"> public class HelloWorld { public static void main(String[] args) { System.out.println("Hello, World"); } } </pre> <p data-bbox="1163 735 1325 816">/bin/java /opt/app.jar /lib/libc</p> <p data-bbox="1213 943 1276 971">layer</p> <p data-bbox="1430 943 1619 971">image index</p> <p data-bbox="1801 943 1871 971">config</p> <pre data-bbox="1430 703 1577 816"> { "manifests": { "platform": { "os": "linux", ... } } } </pre> <pre data-bbox="1724 654 1871 865"> { ... "config": { "Cmd": ["java", "-jar", "app.jar"], ... } } </pre> </div> <p data-bbox="674 1000 1520 1068">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="690 207 1335 267">OCI Image Configuration</h2> <p data-bbox="690 321 1955 483">An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in Layers.</p> <p data-bbox="690 521 1696 553">This section defines the <code>application/vnd.oci.image.config.v1+json</code> media type.</p> <p data-bbox="672 586 1543 656">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 1	Accused Instrumentalities
	<p data-bbox="701 215 789 248">Layer</p> <ul data-bbox="730 289 1957 638" style="list-style-type: none"> • Image filesystems are composed of <i>layers</i>. • Each layer represents a set of filesystem changes in a tar-based layer format, recording files to be added, changed, or deleted relative to its parent layer. • Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer. • Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem. <p data-bbox="701 686 898 719">Image JSON</p> <ul data-bbox="730 760 1957 1109" style="list-style-type: none"> • Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes. • The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers. • This JSON is considered to be immutable, because changing it would change the computed ImageID. • Changing it means creating a new derived image, instead of changing the existing image. <p data-bbox="674 1133 1543 1206">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 1	Accused Instrumentalities
	<ul style="list-style-type: none">• rootfs <i>object</i>, REQUIRED <p>The rootfs key references the layer content addresses used by the image. This makes the image config hash depend on the filesystem hash.</p> <ul style="list-style-type: none">◦ type <i>string</i>, REQUIRED <p>MUST be set to <code>layers</code>. Implementations MUST generate an error if they encounter a unknown value while verifying or unpacking an image.</p> <ul style="list-style-type: none">◦ diff_ids <i>array of strings</i>, REQUIRED <p>An array of layer content hashes (<code>DiffIDs</code>), in order from first to last.</p> <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 2

Claim 2	Accused Instrumentalities
<p>2. A method as defined in claim 1, wherein each container has an execution file associated therewith for starting the one or more applications.</p>	<p>HPE and/or its customer practices, through the Accused Instrumentalities, a method as defined in claim 1, wherein each container has an execution file associated therewith for starting the one or more applications.</p> <p>For example, a container image has an associated image configuration comprising information for starting the one or more applications. This can be an Open Containers Initiative image configuration.</p> <p><i>See, e.g.:</i></p> <p>Open Container Initiative</p> <hr/> <p>Image Format Specification</p> <hr/> <p>This specification defines an OCI Image, consisting of an image manifest, an image index (optional), a set of filesystem layers, and a configuration.</p> <p>The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p> <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p>

Claim 2	Accused Instrumentalities
	<p data-bbox="688 212 869 250">Overview</p> <p data-bbox="688 306 1934 548">At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more filesystem layer changeset archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p> <div data-bbox="709 589 1948 971"> <pre data-bbox="709 727 1020 818"> public class HelloWorld { public static void main(String[] args) { System.out.println("Hello, World"); } } </pre> <p data-bbox="1163 737 1325 818">/bin/java /opt/app.jar /lib/libc</p> <p data-bbox="1213 943 1272 971">layer</p> <p data-bbox="1430 678 1577 818"> { "manifests": { "platform": { "os": "linux", ... } } } </p> <p data-bbox="1472 943 1612 971">image index</p> <p data-bbox="1724 656 1877 867"> { ... "config": { "Cmd": ["java", "-jar", "app.jar"], ... } } </p> <p data-bbox="1801 943 1871 971">config</p> </div> <p data-bbox="674 997 1520 1068"> https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md </p>

Claim 2	Accused Instrumentalities
	<h2 data-bbox="688 207 1335 266">OCI Image Configuration</h2> <p data-bbox="688 321 1955 483">An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in Layers.</p> <p data-bbox="688 521 1696 553">This section defines the <code>application/vnd.oci.image.config.v1+json</code> media type.</p> <p data-bbox="672 586 1541 656">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 2	Accused Instrumentalities
	<ul style="list-style-type: none"> • config object, OPTIONAL The execution parameters which SHOULD be used as a base when running a container using the image. This field can be <code>null</code>, in which case any execution parameters should be specified at creation of the container. <ul style="list-style-type: none"> ◦ Env array of strings, OPTIONAL Entries are in the format of <code>VARNAME=VARVALUE</code>. These values act as defaults and are merged with any specified when creating a container. ◦ Entrypoint array of strings, OPTIONAL A list of arguments to use as the command to execute when the container starts. These values act as defaults and may be replaced by an entrypoint specified when creating a container. ◦ Cmd array of strings, OPTIONAL Default arguments to the entrypoint of the container. These values act as defaults and may be replaced by any specified when creating a container. If an <code>Entrypoint</code> value is not specified, then the first entry of the <code>Cmd</code> array SHOULD be interpreted as the executable to run. <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 6

Claim 6	Accused Instrumentalities
<p>6. A method as defined in claim 2, comprising the step of assigning a unique associated identity to each of a plurality of the containers, wherein the identity includes at least one of IP address, host name, and MAC address.</p>	<p>HPE and/or its customer practices, through the Accused Instrumentalities, a method as defined in claim 2, comprising the step of assigning a unique associated identity to each of a plurality of the containers, wherein the identity includes at least one of IP address, host name, and MAC address.</p> <p>For example, Kubernetes containers have an associated hostname, which in the case of a single-container Pod is the unique identity of that container. For another example, Kubernetes pods have an associated hostname, which is unique. For another example, a networked Kubernetes pod has an assigned IPv4 and/or IPv6 address. For another example, a Docker container has an IP address and a hostname.</p> <p><i>See, e.g.:</i></p> <p>Container information</p> <p>The <i>hostname</i> of a Container is the name of the Pod in which the Container is running. It is available through the <code>hostname</code> command or the <code>gethostname</code> function call in libc.</p> <p>The Pod name and namespace are available as environment variables through the downward API.</p> <p>User defined environment variables from the Pod definition are also available to the Container, as are any environment variables specified statically in the container image.</p> <p>https://kubernetes.io/docs/concepts/containers/container-environment/</p>

Claim 6	Accused Instrumentalities
	<h2 data-bbox="688 207 1289 256">IP address and hostname</h2> <p data-bbox="688 302 1948 423">By default, the container gets an IP address for every Docker network it attaches to. A container receives an IP address out of the IP subnet of the network. The Docker daemon performs dynamic subnetting and IP address allocation for containers. Each network also has a default subnet mask and gateway.</p> <p data-bbox="688 475 1948 643">You can connect a running container to multiple networks, either by passing the <code>--network</code> flag multiple times when creating the container, or using the <code>docker network connect</code> command for already running containers. In both cases, you can use the <code>--ip</code> or <code>--ip6</code> flags to specify the container's IP address on that particular network.</p> <p data-bbox="688 695 1948 816">In the same way, a container's hostname defaults to be the container's ID in Docker. You can override the hostname using <code>--hostname</code>. When connecting to an existing network using <code>docker network connect</code>, you can use the <code>--alias</code> flag to specify an additional network alias for the container on that network.</p> <p data-bbox="672 854 1100 881">https://docs.docker.com/network/</p>

Claim 9

Claim 9	Accused Instrumentalities
<p>9. A method as defined in claim 2, wherein server information related to hardware resource usage including at least one of CPU memory, network bandwidth, and disk allocation is associated with at least some of the containers prior to the applications within the containers being executed.</p>	<p>HPE and/or its customer practices, through the Accused Instrumentalities, a method as defined in claim 2, wherein server information related to hardware resource usage including at least one of CPU memory, network bandwidth, and disk allocation is associated with at least some of the containers prior to the applications within the containers being executed.</p> <p>For example, Kubernetes tracks and limits resource usage, including CPU and memory resources. For another example, Docker tracks and limits resource usage, including CPU and memory resources.</p> <p>See, e.g.:</p> <p><u>Resource Management for Pods and Containers</u></p> <p>When you specify a <u>Pod</u>, you can optionally specify how much of each resource a <u>container</u> needs. The most common resources to specify are CPU and memory (RAM); there are others.</p> <p>When you specify the resource <i>request</i> for containers in a Pod, the <u>kube-scheduler</u> uses this information to decide which node to place the Pod on. When you specify a resource <i>limit</i> for a container, the <u>kubelet</u> enforces those limits so that the running container is not allowed to use more of that resource than the limit you set. The kubelet also reserves at least the <i>request</i> amount of that system resource specifically for that container to use.</p> <p>Requests and limits</p> <p>If the node where a Pod is running has enough of a resource available, it's possible (and allowed) for a container to use more resource than its request for that resource specifies. However, a container is not allowed to use more than its resource limit.</p>

Claim 9	Accused Instrumentalities
	<p>For example, if you set a memory request of 256 MiB for a container, and that container is in a Pod scheduled to a Node with 8GiB of memory and no other Pods, then the container can try to use more RAM.</p> <p>If you set a memory limit of 4GiB for that container, the kubelet (and <u>container runtime</u>) enforce the limit. The runtime prevents the container from using more than the configured resource limit. For example: when a process in the container tries to consume more than the allowed amount of memory, the system kernel terminates the process that attempted the allocation, with an out of memory (OOM) error.</p> <p>Limits can be implemented either reactively (the system intervenes once it sees a violation) or by enforcement (the system prevents the container from ever exceeding the limit). Different runtimes can have different ways to implement the same restrictions.</p> <p>https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</p> <p>Runtime options with Memory, CPUs, and GPUs</p> <p>By default, a container has no resource constraints and can use as much of a given resource as the host's kernel scheduler allows. Docker provides ways to control how much memory, or CPU a container can use, setting runtime configuration flags of the <code>docker run</code> command. This section provides details on when you should set such limits and the possible implications of setting them.</p> <p>Limit a container's access to memory</p> <p>Docker can enforce hard or soft memory limits.</p> <ul style="list-style-type: none"> • Hard limits lets the container use no more than a fixed amount of memory.

Claim 9	Accused Instrumentalities
	<ul style="list-style-type: none"> Soft limits lets the container use as much memory as it needs unless certain conditions are met, such as when the kernel detects low memory or contention on the host machine. <p>https://docs.docker.com/config/containers/resource_constraints/</p>

Claim 10

Claim 10	Accused Instrumentalities
<p>10. A method as defined in claim 2, wherein in operation when an application residing within a container is executed, said application has no access to system files or applications in other containers or to system files within the operating system during execution thereof.</p>	<p>HPE and/or its customer practices, through the Accused Instrumentalities, a method as defined in claim 2, wherein in operation when an application residing within a container is executed, said application has no access to system files or applications in other containers or to system files within the operating system during execution thereof.</p> <p><i>See, e.g.:</i></p> <p>Docker container: A <i>Docker container</i> is a lightweight, standalone, executable software package that runs specific services. This software package includes code, runtime, system libraries, configurations, etc. that run as an isolated process in user space. A Docker container is typically used to deploy scalable and repeatable <i>microservices</i>. HPE Ezmeral Runtime Enterprise contains innovations around storage, networking, and security to utilize Docker containers as lightweight virtual machines to run Big Data and analytics applications.</p> <p>https://support.hpe.com/hpesc/public/docDisplay?docId=a00ecp54hen_us&docLocale=en_US&page=home/about-hpe-ezmeral-container-pl/GEN_Definitions.html</p> <p>Kubernetes namespaces have the following uses:</p> <ul style="list-style-type: none"> Isolation: Teams, projects, and customers exist in their own environment within a cluster, and do not impact each other's work. <p>https://support.hpe.com/hpesc/public/docDisplay?docId=a00ecp55hen_us&docLocale=en_US&page=reference/universal-concepts/Namespace.html</p> <p>Because each application container creates an isolated environment for its application, the resources allocated to it are the entire machine. Other copies of the same container are "unaware" of each other.</p> <p>https://developer.hpe.com/blog/kubernetes-application-containers-managing-containers-and-cluster-resour/</p>

